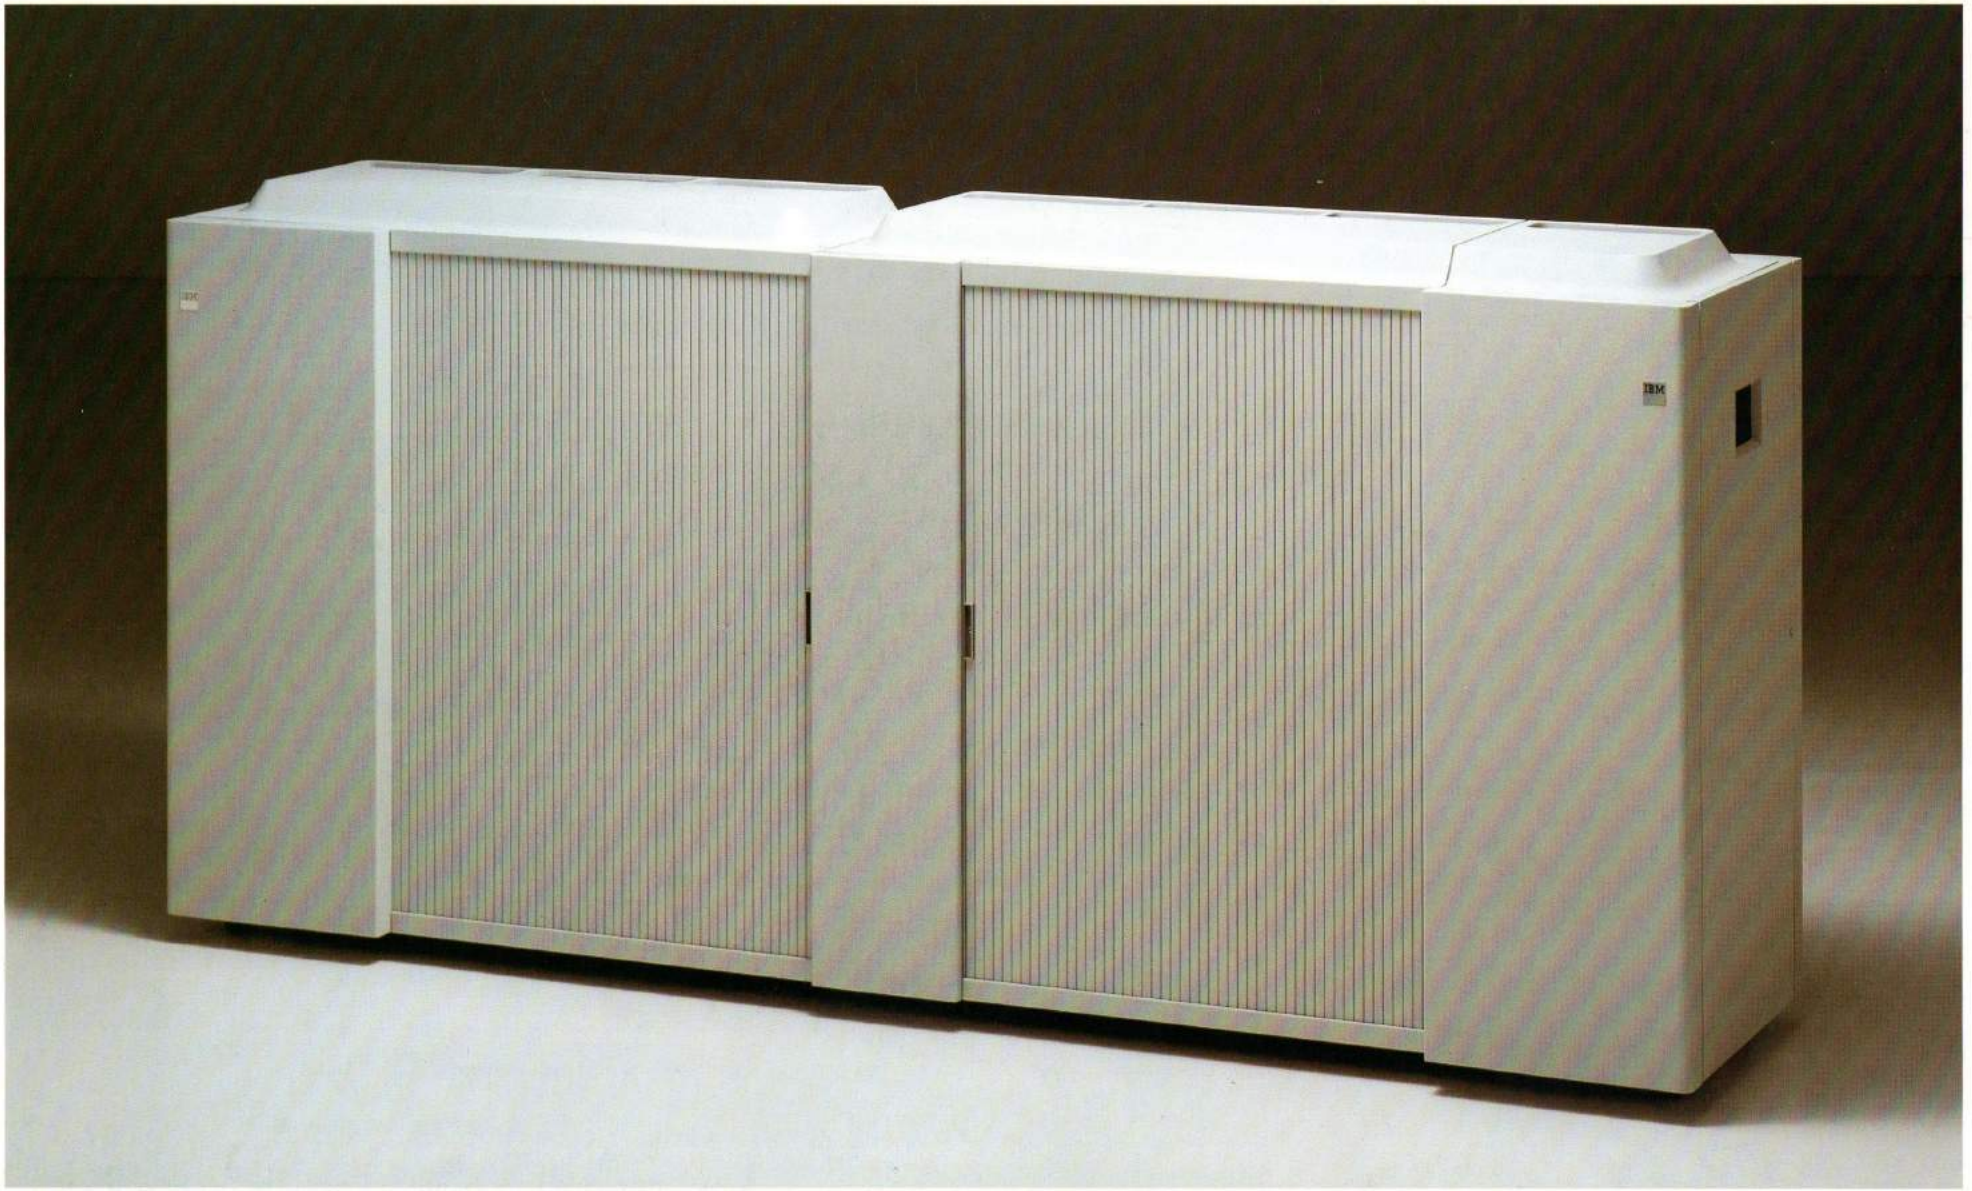

IBM

IBM 3081 Processor Unit Model D

Covering the chips and bolted to the substrate is a metal cap. Its primary function is cooling, via the sealed-in helium gas and springloaded pistons, which are in contact with every chip on the substrate. The cap and helium also have a very important secondary role in protecting the chips against physical damage and chemical contamination.

The more efficient, controlled cooling obtained with the TCM design has meant that storage array chips can be intermixed more freely with logic chips. The greater level of integration achieved has allowed a higher degree of functional design, bringing advantages in performance, testing and diagnostics.

New Multilayered Boards For TCMs

The TCMs plug into nine- or six-position boards consisting of 20 power and signal layers. This uses a design similar to that of the TCM substrate to contain, for example, approximately 3500 wires in the six signal layers.

A new connector design allows the TCMs to be plugged into and out of boards easily. It allows a TCM first to be aligned with its board location, then gradually cammed into position so that all 1800 pins of the TCM are brought into good physical contact with those on the board. A repeatable, low-resistance contact is obtained by a design in which each board pin adjusts itself to the module pin as the TCM is cammed into position. Signals and power to and from the boards are routed by specially designed cables that plug into the board sides.

Very Few Cables

A directly visible effect of the TCM/board design is the neat and almost cableless interior. In the 3081, only 6 percent of connections are made by cables and wires. Thus, with far fewer wires to be routed and shielded properly to avoid noise problems, the possibility of random system malfunctions is significantly reduced.

Extensive Use Of Microcode

The 3081 contains multiple, distributed control stores and logic, which interact with each other asynchronously. This extensive use of microcode is a natural complement to the dense circuitry of the TCM as it can permit functional changes without impacting hardware and provide the basis for operational enhancements.

Other highlights are:

– All instructions are under microcode control to give them greater flexibility. Thus, this particular benefit, traditionally only found in smaller processors, has been extended upwards to the 3081.

– A portion of the control store from which processor microcode executes is dynamically loaded. This is designed to eliminate traditional microcode capacity constraints and choices between mutually exclusive capabilities. Just as virtual storage removed real storage constraints for application programming, dynamically loaded microcode has removed capacity constraints for microcode.

A Simplified Implementation

The 3081 is functionally divided into:
– two central processors, system controller, central storage and external data controller (channels). A second functional unit is the separate 3082 Processor Controller.

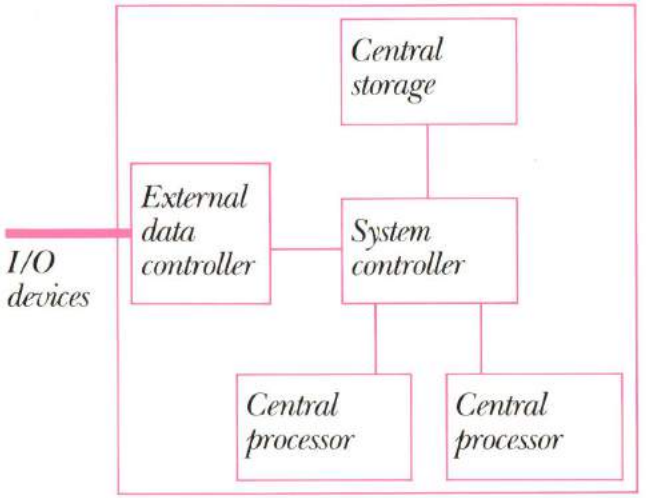

IBM 3081 Processor Unit Model D

The IBM 3081 Processor Unit Model D is a large capacity, high-performance processor.

The model features:

– a dyadic processor consisting of two integrated central processors, sharing the same central storage and operating under a single control program

– the Thermal Conduction Module (TCM), which is a technological evolution in interconnecting, powering, and cooling high-circuit density silicon chips.

System/370 Extended Architecture

The IBM 3081 Processor Unit Model D will operate under the current IBM System/370 mode of operation, as well as IBM's System/370 Extended Architecture. The latter is specifically designed for IBM's very large processors to allow them to reach their full potential in certain environments.

The above advances result in a processor with an instruction execution rate up to 2.1 times faster than a similarly configured IBM 3033-U Processor, when operating in an MVS/SP Version 1.3 environment.

The enhanced performance of the processor is matched by up to 24 integrated channels and a central storage capacity of 16, 24, or 32 Mb.

The channels can be all block multiplexer channels capable of supporting data streaming and individual data transfer rates of 3 Mb/second, or can include up to four byte multiplexer channels.

The 16, 24, and 32 Mb central storage can be directly provided by Models D16, D24, and D32 respectively. Alternatively, you can initially install a Model D16, then upgrade it, as and when required.

These are just a few, among many, of the innovative design features to be found in the IBM 3081 Processor Unit Model D. Some are due to hardware improvements, some the result of extensive microcode development. Together, they combine to provide you with the increased throughput capability you need to support the extensive online applications now necessary in the challenging world of modern business.

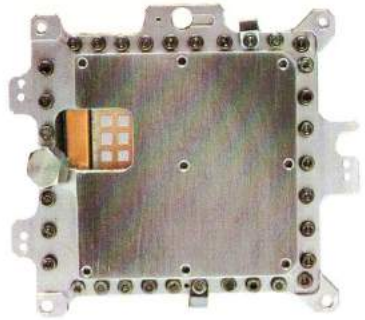

The Thermal Conduction Module

Each Thermal Conduction Module is a helium gas filled assembly, approximately 12.5 cm (5 in.) square. A logic function TCM houses up to 100 silicon chips mounted directly on a ceramic substrate; a storage function TCM houses up to 118 chips. The total number of circuits in a TCM can be more than 45,000, the equivalent of an IBM System/370 Model 148.

The substrate consists of 33 ceramic layers, each one covered by a maze-like pattern of lines deposited by molybdenum metallisation and electrically interconnected. This technique is designed to:

– improve reliability, by eliminating many thousands of pin-to-wire connections. The sixteen layers that provide signal paths, for example, collectively contain the equivalent of 6000 wires

– increase performance, by packing chips more closely together and by significantly reducing the distances that signals have to travel.

Central Processors

As a dyadic processor, the IBM 3081 Model D integrates the two central processors within a single unit operating under a single control program.

Each processor:

– can access the central storage

– has its own 32 Kb buffer storage, which is transparent to programs and helps to improve effective central storage access time

– can transfer data between central storage and the buffers in 128 byte blocks.

The store in buffer design further improves buffer efficiency by reducing traffic to central storage. (Channel interaction with central storage is direct and bypasses buffer storage.)

The central processors translate virtual storage addresses by means of a Dynamic Address Translation (DAT) facility.

To reduce the time for the translation of virtual addresses, the DAT facility makes use of a Translation Lookaside Buffer (TLB) that can store up to 128 addresses.

System/370 Extended Architecture

Under System/370 Extended Architecture central processor performance is aided by:

– offloading part of I/O processing to the External Data Controller

– saving central processor cycles if the channel is busy, or if other busy conditions occur. The software, in effect, does not have to try again.

System Controller

The System Controller is the central switching element for all communications among processor elements. It also controls the transfer of data between central storage and all other elements.

An eight-position queue allows up to eight fetch or store requests to be stacked in various stages of completion. Some overlap of data transfer is possible. For example, a fetch operation might overlap a store operation in certain circumstances.

Among other system-related functions that are part of the controller are time-of-day clock, central storage reconfiguration, and storage protect keys.

Central Storage

The central storage is two-way interleaved, and is available in three capacities:

– 16 Mb (16 777 216 bytes)

– 24 Mb (25 165 824 bytes)

– 32 Mb (33 554 432 bytes)

It has a storage distribution element that contains the logic for fetching or storing doublewords into the data arrays. Each 2K (3081-D16) or 4K (all other 3081 processor models) of storage is protected by one of 15 possible key values to prevent unauthorised access to information in the storage. For store operations, the system controller is notified if a storage violation is attempted, and the data is not stored; the same protection is available for fetch operations.

Error correcting code bits are stored in the data arrays, along with the data. These bits allow the detection and correction of single-bit errors, and the detection of double-bit errors.

External Data Controller

The External Data Controller provides the channel and I/O interfaces for attachment of I/O devices and subsystems. It is microcode controlled allowing efficient operation in both System/370 mode and the System/370 Extended Architecture mode.

The basic number of channels is 16. These can be all block multiplexer channels or can include up to four byte multiplexer channels. You can add eight more block multiplexer channels by installing the Channel Group Additional feature.

Channel addresses are not fixed. They can be assigned at installation time, and can be changed from the system console whenever necessary.

Availability

– channel-provided retry data to enable device-dependent error recovery routines to retry an I/O operation when an error occurs

– automatic access, in the event of a failure of one central processor, to the remaining central processor's channels in 370 mode. In 370-XA mode all channels are accessible to both central processors all the time.

System/370 Extended Architecture– Dynamic Channel Subsystem

IBM's System/370 Extended Architecture introduces some evolutionary changes to the channels which can allow more devices, greater I/O throughput, and improved response times. The Extended Channel Architecture is implemented in the Dynamic Channel Subsystem on the 3081 Model D, and it means that:

– more of the I/O processing can be offloaded to the Dynamic Channel Subsystem

– the Dynamic Channel Subsystem performs I/O path selection and can handle request queuing for all devices. In addition, it can measure I/O requests and queuing times which provides better performance data and can also be used by MVS for I/O scheduling

– there is no central processor affinity between channels; essentially, all the channels are accessible to both processors

– in the event of processor failure, channel set switching and reconfiguration is not required

– dynamic pathing with the appropriate feature on the IBM 3880 provides the potential for reduced response times and greater I/O throughput.

Byte and Block Multiplexer Channels

Byte multiplexer channels can operate in burst or byte mode. When operating in byte mode, they provide for the attachment of low- to medium-speed I/O devices, such as card readers, punches, and terminals.

The performance of byte multiplexer channels is highly dependent on I/O unit interface transition response times, which vary widely from device to device. With the maximum of four byte multiplexer channels installed, the aggregate data transfer rate can range from 45 Kb per second to 450 Kb per second.

The single data path of the channel may be preempted by the I/O control unit for burst mode operation, or may be shared by many I/O devices in byte mode operation. In either case, data transfer is controlled one byte at a time. The sharing of the data path by many devices makes each device appear to the programmer as if it had a data path of its own (a subchannel).

The information that applies to a single subchannel is stored in a unit control word. Each byte multiplexer channel can have 256 subchannels, of which up to eight can be shared and configured to control 8, 16, or 32 devices on one control unit. When a subchannel is shared, the number of unshared subchannels must be reduced by 8, 16, or 32 to correspond to the number of I/O devices on the shared subchannel.

Block multiplexer channels are designed to increase system throughput by increasing the amount of data entering and leaving the system in a given period of time (the effective data transfer rate). Data moves to or from an I/O device one byte at a time, but two doublewords are buffered for communication with central storage. The channels support data streaming and are capable of transferring data at 3 Mb/second.

A single block multiplexer channel can support interleaved, concurrent execution of multiple high-speed I/O operations, in much the same way as byte multiplexer channels can be shared by multiple low-speed devices. Each block multiplexer channel can address up to 256 I/O devices or subchannels and may physically attach up to eight control units.

A block multiplexer channel may operate in selector or block multiplexer mode. In selector mode, the channel remains busy for the entire time the channel program is in operation, whether or not data is being transferred; whereas, in block multiplexer mode, the channel disconnects from the channel program during certain non-data transfer operations. It, therefore, becomes available for an I/O operation on another device, thus enhancing system availability. The mode of operation can be selected under program control.

Modes of Operation

System 370 mode—*System 3081 Model D operates under VM/System Product and MVS/System Product Version 1. The performance under MVS/SP is significantly enhanced by the 3033 Extension Feature and the 3033 Extended Addressing Feature which are included as standard in the 3081 Model D. These provide the following services:*

– cross memory services
– I/O queuing function in the channels
–microcode assist for frequently used control program sequences
–the capability to access up to 32Mb of real storage.

Similarly, for VM users, Virtual Machine Assist and Preferred Machine Assist are included as basic.

When operating in System/370 mode, the integrated dyadic structure of the 3081 offers certain availability benefits.

Should, for example, one central processor fail, its channel set can be accessed by the remaining central processor, thus processing can continue until a suitable time can be allocated for maintenance.

Extended Architecture mode—*The 3081 Model D operates under MVS/XA, which includes MVS/System Product Version 2 which is upward compatible from MVS/SP Version 1. The Extended Architecture provides functional and performance enhancements through:*

– a real address capability up to 2Gb (2 Gigabytes= 2048 Mb)
– a virtual address capability up to 2Gb (2048 Mb)
– a new channel mode of operation, with an identical data format and fully compatible at the control unit and device level with System/370 mode
– and new instructions for all of the above.

The hardware and microcode enhancements found in MVS/SP Version 1 have equivalents in Version 2, thus providing a similar enhanced performance.

When operating in Extended Architecture mode, there are no affinities between the central processors and the channels as in System/370 mode. Instead, either central processor can initiate I/O operations on any device and may also process the I/O completion interrupt from any device. The External Data Controller assumes responsibility for path selection. Failure of one of the central processors will not affect the way in which the 3081 handles I/O requests.

Summary of Characteristics

Standard functions:

– Cycle time of 26 ns.
– 16, 24, or 32 Mb central storage.
– Two central processors integrated as a single, dyadic processor.
– 32 Kb high-speed buffer per central processor.
– Compatibility with System/370 and 303X processors maintained. Basic Control (BC) and Extended Control (EC) modes of operation provided.
– Dynamic Address Translation (DAT) hardware translates virtual storage addresses into real storage address during program execution. Available only in EC mode.
– Channel indirect addressing is used when DAT mode is in effect. Because an I/O buffer may be assigned to noncontiguous real storage areas, channel indirect addressing is required to access the list of the real storage areas (page frames) assigned.
– Channel-provided retry data to enable device-dependent error recovery routines to retry an I/O operation when an error occurs.
– Universal instruction set. Supports binary, decimal, and floating point arithmetic, and contains several general purpose and control instructions.
– Extended precision floating point provides precision of up to 28 hexadecimal digits, equal to up to 34 decimal digits.
– Overlap of instruction fetching with instruction execution.
– Overlap of operand fetching in storage-to-storage instructions.
– High-speed multiply is basic to the implementation of execution element. It incorporates a new algorithm.

– *A monitor function can be used to trace user-defined program events for debugging or gathering statistics.*
– *Program Event Recording can be used to monitor certain events that might occur during program execution (for example, successful branching or alteration of general registers).*
– *Byte (as opposed to halfword, word or doubleword) boundary alignment is permitted for the operands of nonprivileged instructions to reduce the need for padding in records for the purpose of aligning fixed or floating point data.*
– *An interval timer of 3.33 ms is provided for job accounting.*
– *A time-of-day clock that is updated every microsecond is included to provide time-of-day values in conjunction with available programming support. A clock security switch provides an interlock against unauthorised setting or changing of the time-of-day clock.*
– *A clock comparator provides an interrupt when the time-of-day clock reaches a program-specified value.*
– *A processor timer measures elapsed time. It has an interval timing capability equivalent to that of the interval timer, with a resolution equal to that of the time-of-day clock.*

System/370 Mode Only

– *Automatic access, in the event of a failure of a central processor, to its assigned channel set by the remaining central processor.*
– *3033 Extended Addressing, which is designed to increase the addressable area of central storage. More data and programs can remain in central storage at any time.*
– *3033 Extension Feature to offer enhanced performance in an MVS environment.*
– *Virtual Machine Assist to help increase performance for VM users.*
– *Preferred Machine Assist to help increase performance for VM/MVS users.*
– *System/370 Extended Facility is provided to help reduce the time needed to execute certain frequently used supervisor functions of MVS, and increase the efficiency of Dynamic Address Translation.*

System/370 Extended Architecture only

– *A real address capability up to 2Gb (2048 Mb)*
– *A virtual address capability up to 2Gb (2048 Mb)*
– *A new channel mode of operation (Extended Architecture Channels), with an identical data format and fully compatible at the control unit and device level with System/370 mode.*
– *New instructions for the above functions.*

Special Feature

– *Channel Group Additional: provides an additional group of eight channel block multiplexer channels. Requires support from Model 24 of 3082 Processor Controller.*

Physical Size

The dimensions given are indicative only. Service Area dimensions are not included.

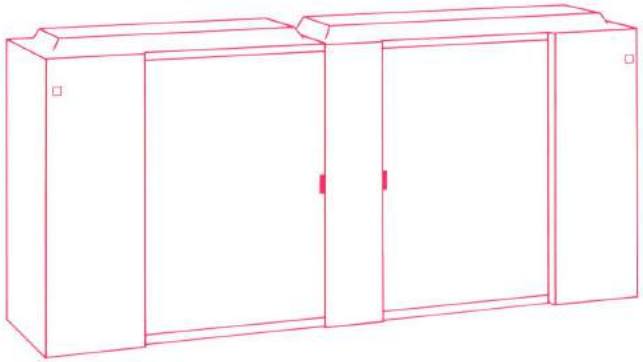

Depth	*Height*	*Length*
813 mm (32 in)	*1778 mm (70 in)*	*3905 mm (154 in)*

The pictures show design models only.

It is possible that this publication may contain references to, or information about, IBM products (machines or programs), programming or services which are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such products, programming or services in your country.

IBM Eurocoordination
Tour Franklin
Cedex 11
92081 Paris, La Défense
France

IBM World Trade
Americas/Far East Corporation
Town of Mount Pleasant
Route 9, North Tarrytown
New York, 10591, USA

Printed in USA

GK 10-6366 - 00 (10/81)

IBM